

Serial Powering for the Tracker Phase-2 Upgrade

Dominik Koukola^{*†}

On behalf of the CMS Collaboration

E-mail: dominik.koukola@cern.ch

The high luminosity phase of the LHC at CERN will bring unprecedented operational conditions for the experiments. To cope with the harsh radiation environment and extreme hit rates in the innermost layers of the detector, both ATLAS and CMS will upgrade their inner tracker pixel detectors. A serial powering distribution scheme is the only viable solution to supply them with the required current levels, within acceptable material budget and power cable losses. This powering scheme is based on the use of Shunt-Low Dropout regulators that will be integrated in the future pixel readout chip designed by the RD53 collaboration. In this paper, this scheme will be explained based on the implementation choices for the future CMS Inner Tracker. Its main system challenges will be discussed and first measurements with serially powered chains of prototype modules will be shown.

The 28th International Workshop on Vertex Detectors - Vertex2019

13-18 October, 2019

Lopud, Croatia

*Speaker.

†Special thanks to A. Nidriche (University of Strasbourg) for her contribution to this work.

1. Introduction

The upgrade of the Large Hadron Collider at CERN for the High-Luminosity (HL-LHC) phase will introduce extreme operating conditions for the experiments. The CMS Inner Tracker (IT) will have to cope with hit rates of up to 3.2 GHz/cm^2 , a L1 trigger rate of 750 kHz as well as a L1 trigger latency of $12.5 \mu\text{s}$. The radiation levels will be reaching unprecedented levels of 1.2 Grad of total ionizing dose and a hadron fluence of $2.3 \times 10^{16} \text{ n}_{\text{eq}} \text{ cm}^{-2}$ for an integrated luminosity of 3000 fb^{-1} . In addition, there are very tight physical constraints on the available space for the services as well as stringent material requirements in order to maintain good tracking performance. Therefore, a new and highly efficient, low mass IT detector is being designed and will be installed in CMS to fully exploit the physics potential of the upgraded machine [1].

The IT detector will be made of ~ 3900 hybrid modules i.e. pixelated silicon sensors bump-bonded to readout chips (two or four). Both thin planar and 3D silicon sensors are being considered with a pixel aspect ratio of $25 \mu\text{m} \times 100 \mu\text{m}$ or $50 \mu\text{m} \times 50 \mu\text{m}$. The pixel readout chip demands a deep sub-micron CMOS technology and an efficient architecture, where groups of pixel channels will share digital resources for buffering, control and data formatting. A large demonstrator pixel readout chip, the RD53A [2], has been designed by the RD53 collaboration (ATLAS and CMS) in 65 nm CMOS technology. The use of this modern high density low power CMOS technology with low supply voltage (1.2 V), demands significant current levels of about 2 A per chip, resulting in a total of 40 kW needed to power the pixel readout chips in the IT. The CMS specific version of the RD53 chip will be submitted in spring 2020.

2. Serial powering scheme




Figure 1: Serial powering scheme with a chain of modules.

The ATLAS and CMS experiments have chosen a serial power distribution scheme to power the pixel modules. It is the only viable solution to supply the pixel readout chips with the required current levels, with acceptable material budget and power cable losses, while also being sufficiently

radiation hard. This powering scheme has never been used in high energy physics on a large scale before and therefore brings new implications compared to classical parallel powering schemes. In the following this scheme will be explained based on the implementation choices for the future CMS IT, while many things are also applicable to the ATLAS pixel detector [3].

In the serial powering scheme, a chain of up to twelve modules is connected in series and powered by a constant supply current as shown in Fig. 1. The current flowing through the chain is used by each module and passed on to the next one. This allows to reduce the number of wires by connecting more modules to a chain, as only one supply line is needed per chain. Unlike classical parallel powering schemes, which are voltage driven, serial powering is current based and thus less sensitive to voltage drops on the supply cables or between modules. Therefore, thinner wires can be used. As the modules are connected in series, each module is exposed to a different voltage potential and therefore has a different local ground, which needs to be electrically isolated from the others.

Within a module, chips are connected in parallel, sharing the incoming current among themselves. A module hosts two or four chips depending on the location in the detector. For the serial powering scheme, on-chip shunt-regulators are needed to generate locally the required voltages for the pixel chip and assure proper current sharing between the chips. The RD53A chip has two integrated, radiation hard, and highly specialized and optimized Shunt-Low Dropout (Shunt-LDO) regulators, one per power domain (analog and digital). As the voltage is generated locally on-chip, no auxiliary on-detector electronics are needed for powering the pixel modules.

3. Shunt-LDO regulator and RD53A readout chip

The Shunt-LDO regulator combines two functionalities: a shunt part that regulates the input voltage and shunts not consumed current to ground, and a LDO part that regulates the output voltage. A simplified model of the Shunt-LDO in the regulation domain is shown in Fig. 2 on the left, where the input behavior is described as a resistor and an offset voltage in series. The resistive input behavior is essential for the current sharing between parallel chips and regulators. The output behavior is described as a simple voltage source like an ideal LDO. The LDO requires the input voltage to be at least 0.2 V higher than the output voltage for stable regulation.




Figure 2: Simplified Shunt-LDO model in the regulation domain (left) and ideal I-V curve of a RD53A chip with its two Shunt-LDO regulators (right).

Both the input and the output parameters are configurable. The input impedance and offset voltage are defined by external resistors that will be placed on the module. The output voltage is derived from an internal reference voltage that is adjustable with chip internal trim bits. While the output voltage is required to be 1.2 V for optimal performance of the analog and digital parts of the pixel chip, the input parameters can be chosen depending on the needs in the system. Together with the supplied input current they define the working point of the Shunt-LDO regulator.

An ideal I-V curve of the RD53A with its two parallel Shunt-LDO regulators is shown in Fig. 2 on the right. The lowest possible working point in that configuration is at an input current of 1 A, where the input voltage is 1.4 V and the output voltages are 1.2 V. At this point and above, the two Shunt-LDOs are in the regulation domain and provide a stable output voltage.

Parameter		Range
Input current	I_{in}	2 A max
Input voltage	V_{in}	$1.4 \text{ V} < V_{in} < 2 \text{ V}$
Load current	I_{load}	$I_{load} \leq I_{in}$, 1 A max
Output voltage	V_{out}	1.2 V typical

Table 1: Operational range of one Shunt-LDO regulator.

For a reliable and sustainable operation it is important that the Shunt-LDO does not exceed its operational range, specified in Tab. 1. Despite the fact that the RD53A chip is only about half the size of the final chip and therefore also only consumes about half the current, the Shunt-LDOs implemented in the RD53A are already designed for the final production chip. Each Shunt-LDO is capable of providing up to 1 A of output current and can take twice as much (2 A) as input current. This is particularly important in case of failures. Operating at a too low input voltage or drawing too much current will result in unstable regulation, while exceeding the current or voltage limits will damage the regulator's circuitry.

4. System challenges

4.1 Implications of a constant supply current

Typically, an electronic circuit is powered by a constant voltage and draws current depending on its activity and configuration. In serial powering, however, a constant current is provided. This current is also the maximum current that can be drawn by the loads. Therefore, it needs to be assured that enough current is provided to cover the maximum expected load current. In case a pixel chip wants to draw more current than provided, the affected Shunt-LDO regulator will fall out of regulation.

In addition to the maximum expected load current, an additional current headroom needs to be supplied to cover unexpected load peaks and account for current sharing mismatches among the chips on a module. Measurements from wafer probing show that a mismatch of up to $\pm 5\%$ of input impedance and offset voltage between chips should be expected. However, there has been a new trimming scheme introduced for the final pixel chip, which aims at lowering this number. The mismatches come largely from the production process of the chip, while the contribution of

the external resistors for the configuration of the Shunt-LDOs will be negligible, as high precision components will be chosen. It is important to properly account for these mismatches to prevent overload conditions due to bad current sharing.

Providing constantly the maximum expected current plus a headroom means that the pixel detector is always operated at full power. In addition with the losses in the LDO part of the regulator, this results in an intrinsic inefficiency of around 30 % that is necessary to supply the pixel detector reliably with the serial powering scheme. However, it causes local hotspots, as the excessive power is dissipated in the Shunt-LDOs which are located in very small areas along the chip periphery. Special care therefore has to be taken to provide proper cooling. Constant improvements are being made to optimize the thermal contact to the highly efficient CO₂-bi-phase cooling for the final detector design.

4.2 Optimization

The choice of the current headroom and subsequently the input current as well as the choice of the working point parameters (input resistance and offset voltage) of the Shunt-LDOs need to be optimized in order to have an efficient and reliable powering. In this optimal configuration the power consumed by the Shunt-LDO regulator is minimized, while sufficient current is provided to the chip to operate properly.

A study has been performed to verify the assumption of a needed headroom of 20 % and to determine possible combinations of input impedance and offset voltage. It is required to keep the Shunt-LDO in its operational range (see Tab. 1) and provide enough current for all combinations of parameter mismatches between regulators. As the design of the final chip is still ongoing, the exact final power consumption is not yet known. The maximum current consumption for each domain has been estimated to be 0.75 A. In combination with 20 % headroom this gives an input current of 0.9 A per Shunt-LDO. Additionally, it is required that the Shunt-LDO can stay in the operational region in case that one chip develops an open and the input current doubles (see next section).




Figure 3: Result of a first optimization study showing the possible configuration in white, that allow for proper operation, taking into account mismatches and voltage limitations.

The result of the study is shown in Fig. 3. The white area represents the possible configuration combinations that meet the above restrictions. Optimally a point in the middle of the area will be chosen to leave some margin. In this case it could be an offset voltage of 1 V and an input impedance of around 0.56Ω . Alternatively, the current headroom could be reduced to a minimum to find the most power efficient working point. Since the final load current for the analog and digital domains is expected to be slightly different, further optimization studies will be performed in future.

4.3 Typical failure cases

In the serial powering scheme foreseen in the CMS IT, up to 48 chips are connected together in one chain. Therefore failure propagation is a big concern and special features have been developed to address it. Typical failure cases are open or shorts of a chip or a single regulator. These cases are depicted in Fig. 4 using the example of a module with two chips. In normal operation, the current is shared between the chips on the module (left). In case of an open failure of one chip (middle), the chip in parallel needs to take the current of the other chip in addition. This will increase the input voltage on the chip due to resistive behavior and therefore also the total power dissipated on the module. Depending on the configured input impedance this can be more or less severe. In the case of a short the current will flow through the short leaving the parallel chip unpowered. Given the system design, only the chips in parallel are affected by a single chip failure, while the remaining modules in the serial powering chain stay operational.




Figure 4: Operational scenarios of a module with two chips: normal (left), open failure on one chip (middle) and short failure on one chip (right).

Since the submission of RD53A many improvements have been made in the Shunt-LDO design concerning the stability, especially during start up. Additionally two new features have been added to improve the reliability of the regulator. An overvoltage protection which limits the maximum input voltage and an overload protection which protects from temporary overload conditions (e.g. by misconfiguration of the chip) by damping the transient propagation to neighbouring chips. Additionally there are decoupling capacitors placed close to the chips on the module to decouple fast transients $O(>MHz)$. These features are sufficient to enhance the reliability of the scheme and limit the impact of failing chips to the other chips in the chain.

5. Recent system tests and developments for CMS

In the past, extensive tests have been done with single RD53A chips and dedicated Shunt-LDO test chips, which demonstrated the robustness and reliability on a chip level [5]. At the same time, CMS specific developments have been ongoing and a first prototype module that hosts four RD53A chips has been designed [6]. A detailed description of the prototype module and its construction

is reported in [7]. With the recent availability of such digital RD53A quad modules (multi-chip assemblies without a silicon sensor) first CMS specific system tests have been performed on structures that closely resemble the electrical characteristics of the final detector.

The CMS IT will be comprised of a barrel part, where the pixel modules are arranged in "ladders", and two endcap parts, where the pixel modules will be arranged in concentric rings on disks [1]. Two mechanical structures have been produced as shown in Fig. 5: A straight structure resembling a ladder of the barrel part and a ring structure resembling a part of a disk. The base of both is made out of aluminium blocks that have a small layer of aluminium nitride for electrical insulation. Cooling pipes are routed inside the blocks underneath the modules to remove the dissipated heat and allow to go to low temperatures with this setup in the future. The modules are fixed mechanically to the aluminium blocks with clamps, no glue or thermal grease is used.




Figure 5: Ladder structure (left) and ring structure (right) with four RD53A quad modules powered in a serial chain each. Cables for readout are not present in the pictures.

On the ladder structure, as shown in Fig. 5 on the left, modules are placed one after the other to form a serial chain. The current is provided through an adapter board (dark blue) and goes from one module to the next. On the last module, the chain is terminated and the current is routed through the modules back to the adapter board.

As in the ring structure the modules are geometrically not placed in a chain, a light aluminium flex circuit has been designed to route the power from one module to the other as shown in Fig. 5 on the right. It can provide low voltage (LV) and high voltage (HV) to four RD53A quad modules in a serial chain and has two 50 μm thick aluminium layers. The top layer routes the HV (in parallel) and the LV (in series) clockwise from module to module. The bottom layer is used as a common return path for both LV and HV. The power to the aluminium flex can either be provided through a connector and another flex cable or with wires directly soldered to the aluminium flex. With an inner radius of 75 mm and an outer radius of 83 mm the dimensions are close to the space available for the innermost ring in the CMS IT. The design can be easily adapted to the final module dimensions for any ring of the disks in the final detector.

I-V curve measurements have been performed on both structures. Different voltage potentials have been measured along the chain. Figure 6 shows the measured input voltages of the two structures and the total net chip voltage, which includes only the voltage on the chips, excluding voltage drops on traces and connectors. As each module has four chips in parallel, only one chip voltage per module has been added, not taking into account the mismatch of trace losses ($\sim 10 \text{ m}\Omega$) between chips on a module. The total input voltage V_{in} of the ladder includes the voltage drops




Figure 6: IV curve for the whole structure of ladder and ring with four quad modules powered in series.

on the connections between the modules and on the return path of the modules. A difference of 125 mΩ to the net voltage or losses of around 30 mΩ per module have been measured, which is in good agreement with simulations (18 mΩ on traces plus 9 mΩ on connector). The total V_{in} of the ring additionally also includes the voltage drops due to the routing of the current between the modules and to the input connectors of the aluminium flex. The difference between ladder and ring is around 85 mΩ. This value is also in good agreement with resistance measurements of the flex without modules (83 mΩ) and with simulations (76 mΩ). Of all the losses a large part (~35 %) is contributed by the connectors.

Significant amount of power is dissipated on the structures. At the typical working point for the RD53A quad module with an input current of 6 A, around 10 W are dissipated on a module and a total of around 45 W on the whole structure. Due to the resistive behavior of the Shunt-LDO regulators the power increases more than linearly with the input current. With an input current of 8 A already around 70 W are dissipated on the structures. This requires proper active cooling to not overheat the modules.




Figure 7: Temperature measurement with four quad modules on the ladder (left) and ring (right) structure powered with an input current of 8 A (dissipating around 70 W) with basic cooling.

The temperature has been measured with an infrared camera for the extreme case of an input current of 8 A and the results are shown in Fig. 7. While on both structures the central parts of the modules are around 40 °C, the pigtailed and the termination connections are hotspots with around 70 °C. Additionally the pigtail of the first module of the barrel structure (left) has a temperature of 90 °C, because the adapter board is not in contact with the cooling plate. A similar hotspot can be seen on the flex cable providing the current to the aluminium flex (right). The temperatures of the modules are within the operational limits of the pixel chip. The next step is to operate at cold temperatures to resemble more detector-like operation.

6. Summary

A serial powering scheme will be used to power the IT pixel detectors of ATLAS and CMS. Despite being a novel powering scheme that has never been used in high energy physics on a large scale, it is by now well established in the communities of ATLAS, CMS and RD53. Extensive tests and simulations have been performed with single RD53A chips and dedicated Shunt-LDO test chips, which have demonstrated robust performance and reliability on a chip level. Special studies have been done to address all the points of concern such as failure propagation and optimization for reliable yet efficient operation. With the recent availability of digital RD53A quad modules, first CMS specific system tests have been performed with structures that closely resemble the electrical characteristics of the final detector, showing promising results.

Many future challenges are still ahead. Large scale system tests with configurations close to the final detector, the validation of the proper powering behavior of the final CMS prototype pixel chip and the optimization of the serial powering parameters (working point, current headroom and cooling) for the CMS IT will be the key topics.

References

- [1] CMS Collaboration, *The Phase-2 Upgrade of the CMS Tracker*, Technical Design Report, CERN-LHCC-2017-009, CMS-TDR-014, <https://cds.cern.ch/record/2272264>
- [2] M. Garcia-Scriveres, *The RD53A integrated circuit*, Manual, CERN-RD53-PUB-17-001, <https://cds.cern.ch/record/2287593>
- [3] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*, Technical Design Report, CERN-LHCC-2017-021, ATLAS-TDR-030, <https://cds.cern.ch/record/2285585>
- [4] M. Karagounis, *An integrated Shunt-LDO regulator for serial powered systems*, Proceedings of ESSCIRC (2009), DOI: 10.1109/ESSCIRC.2009.5325974
- [5] A. Pradas, *RD53A chip susceptibility to electromagnetic conducted noise*, Proceedings of TWEPP 2019, Submitted for publication
- [6] M. Backhaus, *The Upgrade of the CMS Inner Tracker for HL-LHC*, 2019 JINST 14 C10017, DOI: 10.1088/1748-0221/14/10/c10017
- [7] B. Ristic, *Prototype module construction for the high luminosity upgrade of the CMS pixel detector*, This conference